PCI-SIG ENGINEERING CHANGE NOTICE

TITLE:	Root Complex Event Collector Bus Number Association
DATE:	9/10/2018 minor corrections 9/13/2018
AFFECTED DOCUMENTS:	PCI Express Base Specification, Revision 4.0
SPONSOR:	Intel Corporation

Summary of Functional Changes:

This ECN enhances Root Complex Event Collectors (RCECs) to allow them to be associated with Devices located on additional Bus numbers.

Benefits as a Result of the Changes:

This ECN removes the limitation that an RCEC can only be associated with Endpoints on the same Bus, allowing a single RCECs to cover multiple Buses in a Root Complex.

This simplifies the placement of RCECs in Root complexes where the logical placement of RCiEPs may change based on Root Complex formation, such as RCiEPs that support SR-IOV.

Assessment of the Impact:

Optional normative behavior. No impact if the feature is not implemented.

Analysis of the Hardware Implications:

Optional normative behavior. No impact if the feature is not implemented. Only impacts RCECs.

Analysis of the Software Implications:

Optional normative behavior. No impact if the feature is not implemented. New System Software will be required in cases where the association between RCiEPs and RCECs is determined. Existing software that does not evaluate the Capability Version field of the RCEC Endpoint Association Extended Capability will not associate the RCEC with RCiEPs on additional Buses. No impact if System Software does not attempt to associate RCiEPs and RCECs with this Capability.

Analysis of the C&I Test Implications:

No new C&I tests are envisioned. The usual register bit read testing should be added.

Make the following changes to Section 1.3.2.3:

1.3.2.3 Root Complex Integrated Endpoint Rules

- A Root Complex Integrated Endpoint (RCiEP) is implemented on internal logic of Root Complexes that contains the Root Ports.
- An RCiEP must be a Function with a Type 00h Configuration Space header.
- An RCiEP must support Configuration Requests as a Completer.
- An RCiEP must not require I/O resources claimed through BAR(s).
- An RCiEP must not generate I/O Requests.
- An RCiEP must not support Locked Requests as a Completer or generate them as a Requester. PCI Express-compliant software drivers and applications must be written to prevent the use of lock semantics when accessing an RCiEP .
- An RCiEP operating as the Requester of a Memory Transaction is required to be capable of generating addresses equal to or greater than the Host is capable of handling as a Completer.
- An RCiEP is required to support MSI or MSI-X or both if an interrupt resource is requested. If MSI is implemented, an RCiEP is permitted to support either the 32-bit or 64-bit Message Address version of the MSI Capability structure.
- An RCiEP is permitted to support 32-bit addressing for Base Address Registers that request memory resources.
- An RCiEP must not implement Link Capabilities, Link Status, Link Control, Link Capabilities 2, Link Status 2, and Link Control 2 registers in the PCI Express Extended Capability. ~~See Section 7.5.2 for more details.~~
- ~~An RCiEP must signal PME and error conditions through the same mechanisms used on PCI systems. If a Root Complex Event Collector is implemented, an RCiEP may optionally~~ If an RCiEP is associated with an optional Root Complex Event Collector it must signal PME and error conditions through ~~a~~ the Root Complex Event Collector. ~~In this case, an~~
- An RCiEP must not be associated with ~~no~~ more than one Root Complex Event Collector.
- An RCiEP does not implement Active State Power Management.
- An RCiEP may not be hot-plugged independent of the Root Complex as a whole.
- An RCiEP must not appear in any of the hierarchy domains exposed by the Root Complex.
- An RCiEP must not appear in Switches.

Make the following changes to Section 1.3.4:

1.3.4 Root Complex Event Collector

- A Root Complex Event Collector provides support for terminating error and PME messages from RCiEPs .
- A Root Complex Event Collector must follow all rules for an RCiEP.
- A Root Complex Event Collector is not required to decode any memory or I/O resources.

- A Root Complex Event Collector is identified by its Device/Port Type value (see section 7.5.3.2).
- A Root Complex Event Collector has the Base Class 08h, SubClass 07h and Programming Interface 00h.
- A Root Complex Event Collector resides on ~~the same Logical Bus as the RCiEPs it supports.~~ a Bus in the Root Complex. Multiple Root Complex Event Collectors are permitted to reside on a single ~~Logical~~ Bus.
- A Root Complex Event Collector explicitly declares supported RCiEPs through the Root Complex Event Collector Endpoint Association Capability.

Section 1.7 Error Listing and Rules describes all the errors and how the hardware is required to respond when the error is detected.

Make the following changes to Section 6.2.3.2:

6.2.3.2 Error Messages

Error Messages are sent to the Root Complex for reporting the detection of errors according to the severity of the error.

Error messages that originate from PCI Express or Legacy Endpoints are sent to corresponding Root Ports. Errors that originate from a Root Port itself are reported through the same Root Port.

If ~~a~~ an optional Root Complex Event Collector is implemented, errors that originate from ~~an RCiEP may optionally be~~ RCiEPs are sent to the corresponding Root Complex Event Collector. Errors that originate from ~~an RCiEP are reported in~~ a Root Complex Event Collector ~~residing on~~ itself are reported through the same ~~Logical Bus as the RCiEP.~~ Root Complex Event Collector. The Root Complex Event Collector must ~~explicitly~~ declare supported RCiEPs as part of its capabilities; each RCiEP must be associated with no more than one Root Complex Event Collector.

When multiple errors of the same severity are detected, the corresponding error Messages with the same Requester ID may be merged for different errors of the same severity. At least one error Message must be sent for detected errors of each severity level. Note, however, that the detection of a given error in some cases will preclude the reporting of certain errors. Refer to Section 6.2.3.2.3 Error Pollution. Also note special rules in Section 1.4 Error Logging regarding non-Function-specific errors in Multi-Function Devices.

…

Make the following changes to Section 6.2.4.1.1:

6.2.4.1.1 Error Source Identification

In addition to the above logging, a Root Port or Root Complex Event Collector that supports the Advanced Error Reporting Capability is required to implement the Error Source Identification

register, which records the Requester ID of the first ERR_NONFATAL/ERR_FATAL (uncorrectable errors) and ERR_COR (correctable errors) Messages received by the Root Port or Root Complex Event Collector. System software written to support Advanced Error Reporting can use the Root Error Status register to determine which fields hold valid information.

If ~~a Root Complex Event Collector is implemented, errors from~~ an RCiEP ~~may optionally be reported in a Root Complex Event Collector residing on the same Logical Bus as the RCiEP. The Root Complex Event Collector must explicitly declare supported RCiEPs as part of its capabilities. Each RCiEP must be~~ is associated with ~~no more than one~~ a Root Complex Event Collector, the RCiEP must report its errors through that Root Complex Event Collector.

…

Make the following changes to Section 5.3.3.1:

…

Power management software may transition a Hierarchy into a low power state, and transition the Upstream Links of these devices into the non-communicating L2 state. The PCI Express PME generation mechanism is, therefore, broken into two components:

- Waking a non-communicating Hierarchy (wakeup). This step is required only if the Upstream Link of the device originating the PME is in the non-communicating L2 state, since in that state the device cannot send a PM_PME Message Upstream.
- Sending a PM_PME Message to the root of the Hierarchy.

PME indications that originate from PCI Express Endpoints or PCI Express Legacy Endpoints are propagated to the Root Complex in the form of TLP messages. PM_PME Messages identify the requesting agent within the Hierarchy (via the Requester ID of the PME Message header). Explicit identification within the PM_PME Message is intended to facilitate quicker PME service routine response, and hence shorter resume time.

If an RCiEP is associated with a Root Complex Event Collector ~~is implemented~~, any PME indications that originate from that ~~a Root Complex Integrated Endpoint (~~RCiEP must~~) may optionally~~ be reported ~~in a~~ by that Root Complex Event Collector. ~~residing on the same Logical Bus as the RCiEP~~. ~~The Root Complex Event Collector must explicitly declare supported RCiEPs as part of its capabilities; each RCiEP must be associated with no more than one Root Complex Event Collector. Root Complex Event Collectors explicitly identify the logical location of the requesting agent to facilitate quicker PME service routine response.~~

PME indications that originate from a Root Port itself are reported through the same Root Port.

…

Make the following changes to Section 7.9.10:

7.9.10 Root Complex Event Collector Endpoint Association Extended Capability

The Root Complex Event Collector Endpoint Association Extended Capability is implemented by Root Complex Event Collectors. It declares the RCiEPs supported by the Root Complex Event Collector ~~on the same Logical Bus on which the Root Complex Event Collector is located~~. A Root Complex Event Collector must implement the Root Complex Event Collector Endpoint Association Extended Capability; no other PCI Express ~~d~~Device Function is permitted to implement this Capability.

The Root Complex Event Collector Endpoint Association Extended Capability, as shown in Figure 7-203, consists of the PCI Express Extended Capability header followed by a DWORD bitmap enumerating RCiEPs on the same Bus, and optionally an additional range of Bus Numbers that may contain RCiEPs associated with the Root Complex Event Collector. Functions other than RCiEPs (e.g. Root Ports) contained in the range described by this Capability are not associated with this RCEC.

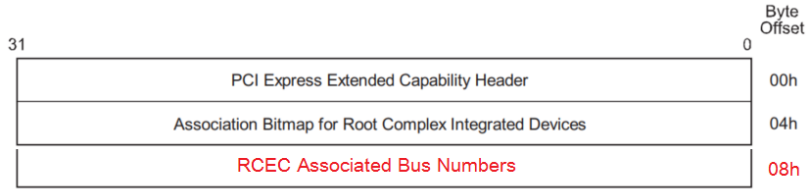

Figure 7-203: Root Complex Event Collector Endpoint Association Extended Capability

Make the following change to Table 7-164:

Table 7-164 Root Complex Event Collector Endpoint Association Extended Capability Header

Bit Location	Register Description	Attributes
15:0	***PCI Express Extended Capability ID*** - This field is a PCI-SIG defined ID number that indicates the nature and format of the Extended Capability. The Extended Capability ID for the Root Complex Event Collector Endpoint Association Extended Capability is 0007h.	RO
19:16	***Capability Version*** - This field is a PCI-SIG defined version number that indicates the version of the Capability structure present. Must be 2h if the Extended Capability contains the RCEC Associated Bus Numbers register (see Section 7.9.10.3). Must be 1h otherwise. ~~for this version of the specification.~~	RO

Table 7-164 Root Complex Event Collector Endpoint Association Extended Capability Header

Bit Location	Register Description	Attributes
31:20	***Next Capability Offset*** - This field contains the offset to the next PCI Express Capability structure or 000h if no other items exist in the linked list of Capabilities. For Extended Capabilities implemented in Configuration Space, this offset is relative to the beginning of PCI-compatible Configuration Space and thus must always be either 000h (for terminating list of Capabilities) or greater than 0FFh. The bottom 2 bits of this offset are Reserved and must be implemented as 00b although software must mask them to allow for future uses of these bits.	RO

Make the following changes to Section 7.9.10.2:

7.9.10.2 Association Bitmap for RCiEPs (Offset 04h)

The Association Bitmap for RCiEPs is a read-only register that sets the bits corresponding to the Device Numbers of RCiEPs ~~supported by~~ associated with the Root Complex Event Collector on the same ~~Logical~~ Bus Number as the Event Collector itself. The bit corresponding to the Device Number of the Root Complex Event Collector must always be Set.

Add Section 7.9.10.3 as follows:

7.9.10.3 RCEC Associated Bus Numbers (Offset 08h)

The RCEC Associated Bus Numbers is a read-only register that indicates an additional range of Bus Numbers containing RCiEPs associated with this Root Complex Event Collector. It is permitted for Functions other than RCiEPs, including Root Ports, to appear within the Association Bus Range. Only RCiEPs in the range are associated with this Root Complex Event Collector. This register is present if the Capability Version is 2h or greater.

This register does not indicate association between an Event Collector and any Virtual Functions within the Association Bus Range (see Section 9.2.1.2). This register does not indicate association between an Event Collector and any Function on the same Bus Number as the Event Collector itself, however it is permitted for the Association Bus Range to include the Bus Number of the Root Complex Event Collector.

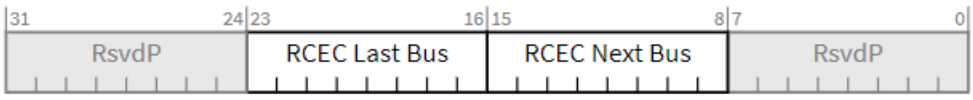

Figure 7-204: RCEC Associated Bus Numbers

Table 7-165: RCEC Associated Bus Numbers

Bit Location	Register Description	Attributes
15:8	***RCEC Next Bus*** **-** This field contains the lowest additional bus number containing RCiEPs associated with this Root Complex Event Collector. If all of the Devices associated with this Root Complex Event Collector are on the same bus as the Event Collector, then this field must be set to FFh.	HwInit
23:16	***RCEC Last Bus*** **-** This field contains the highest additional bus number containing RCiEPs associated with this Root Complex Event Collector. If all of the Devices associated with this Root Complex Event Collector are on the same bus as the Event Collector, then this field must be set to 00h.	HwInit

IMPLEMENTATION NOTE: RCEC Associated Bus Number Compatibility with Legacy Software

Legacy software may not support the use of the RCEC Associated Bus Number register as a mechanism to associate Devices with a RCEC. Such software may see events in the RCEC from Devices on different bus numbers that it does not consider to be associated with the Root Complex Event Collector. System Software is strongly encouraged to report all events seen on the Root Complex Event Collector, regardless of whether or not it can determine association.

Make the following changes to section 9.2.1.2:

9.2.1.2 VF Discovery

...

VFs of an SR-IOV RCiEP Device are associated with the same Root Complex Event Collector (if any) as their PF. Such VFs are not reported in the Root Complex Event Collector Endpoint Association Extended Capability of the Root Complex Event Collector.

...